

614898
5p.

ADVANCES IN ROBOTIC, HUMAN, AND AUTONOMOUS SYSTEMS FOR MISSIONS OF SPACE EXPLORATION

By

Anthony R. Gross¹, Geoffrey A. Briggs¹, Brian J. Glass¹, Liam Pedersen², David M. Kortenkamp³, David S. Wettergreen⁴, I. Nourbakhsh⁴, and Daniel J. Clancy¹

¹NASA Ames Research Center, Moffett Field, CA, USA

²QSS Group, Inc., at NASA Ames Research Center, Moffett Field, CA, USA

³Metrica at Johnson Space Center, Houston, TX, USA

⁴Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

Space exploration missions are evolving toward more complex architectures involving more capable robotic systems, new levels of human and robotic interaction, and increasingly autonomous systems. How this evolving mix of advanced capabilities will be utilized in the design of new missions is a subject of much current interest. Cost and risk constraints also play a key role in the development of new missions, resulting in a complex interplay of a broad range of factors in the mission development and planning of new missions. This paper will discuss how human, robotic, and autonomous systems could be used in advanced space exploration missions. In particular, a recently completed survey of the state of the art and the potential future of robotic systems, as well as new experiments utilizing human and robotic approaches will be described. Finally, there will be a discussion of how best to utilize these various approaches for meeting space exploration goals.

Introduction

Robots have had a role in space exploration from the beginning of humankind's space activities. The Soviet Lunakhod Rover was teleoperated on the surface of the Moon in 1970. More recently, the surface of Mars was

explored by Sojourner, and Remote Manipulator Systems have helped construct the international space station. However, robots have not lived up to the promises initially articulated in such venues as science fiction stories, movies, and TV shows. In particular, current robots lack the reasoning abilities necessary to deal with novel situations and the dexterity to perform human-like manipulation tasks.

With exponentially increasing capabilities of computer hardware and software, including networks and communication systems, a new balance of work is being developed between humans and machines. This new balance holds the promise of greatly increased space exploration capability, along with dramatically reduced design, development, test, and operating costs. New information technologies, which take advantage of knowledge-based software and high performance computer systems, will enable the development of a new generation of design and development tools, schedulers, and vehicle and system health monitoring capabilities. Such tools will provide a degree of machine intelligence and associated autonomy which has previously been unavailable to the mission and spacecraft designer and to the system operator. These capabilities are critical as we look toward future exploration of our solar system, due to both the requirements levied by these missions as well as the

Copyright ©2002 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

Presented at the 53rd International Astronautical Congress, Houston, Texas, USA, October 10-19, 2002.

* Anthony R. Gross is Associate Director of the Information Sciences and Technology Directorate; Geoffrey A. Briggs is Director of the Center for Mars Exploration (CMEX); Brian J. Glass is; Liam Pedersen is; David M. Kortenkamp is; David S. Wettergreen is; I. Nourbakhsh is; Daniel J. Clancy is Chief of the Computational Sciences Division.

budgetary constraints that limit our ability to monitor and control these missions using a standing army of ground-based controllers. At the same time new intelligence and mobility capabilities being made available to robotic systems indicate that now teams of humans and intelligent robots could provide an entirely new level of space exploration capability, greatly extending the reach and safety of landed missions. The suite of possible exploration modalities is now such that new types of exploration missions can be planned with increased confidence of successful implementation and success.

In the next three sections we will look at robotic systems, robotic and human systems working together, and autonomous systems, respectively, to arrive at a comprehensive view of a complete spectrum of new exploration modalities.

Space Robotics

To provide mission designers with appropriate expectations for the roles that robots might play in the next ten to twenty years, NASA has recently completed a comprehensive survey of robotic systems for space exploration [1], including present (2002) and future state of the art in space robotics. The following discussion presents a brief summary of the findings of that report.

Space robotic functionalities are required to support two broad mission classes: planetary surface exploration and in-space operations. The former focuses on robotic mobility, science perception, instrument placement and sample manipulation. The latter focuses on robotic assembly, inspection and repair. In both classes the report also looks at those functionalities unique to human-robot teaming. Figure 1 depicts the various functionalities.

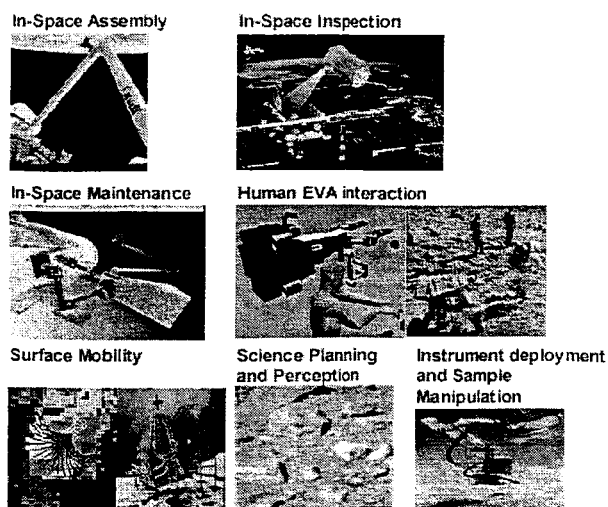


Figure 1. Space Robotic Functionalities

In-Space Assembly

Current in-space capabilities for robotic assembly consist of the space shuttle and space station remote manipulator systems (RMS). These teleoperated robots can move large components and mate those components under careful human teleoperation and supervision. Ground testbeds have demonstrated autonomous transporting and mating of large components (e.g., CMU's Skyworker and NASA Langley's Automated Telescope Assembly). Other ground testbeds have demonstrated teleoperated robots doing fine assembly such as mating connectors (e.g., NASA JSC's Robonaut and University of Maryland's Ranger). In ten years, we expect robots to perform delicate assembly tasks autonomously and even approach the dexterity of a space-suited human. With intense effort, robotic assembly of complicated structures in space is possible, but only under constant supervision and guidance (including occasional teleoperation) from space and ground-based humans. Complex robotic assembly with limited human supervision will require breakthrough technologies.

In-Space Inspection

Currently there are no inspection robots in operation in space. A test of a free-flying camera, AERCam Sprint, was conducted during STS-87 in 1997. This robot was purely teleoperated. In ten years autonomous robotic inspection of some exterior surfaces is feasible. Limited autonomous screening of the sensor data is likely. With intense effort, a robot can autonomously inspect most exterior surfaces and detect anomalies.

In-Space Maintenance

The shuttle and station remote manipulator systems can only move large objects, but do not perform sophisticated maintenance. Several in-space experiments have been performed to demonstrate teleoperated robots doing maintenance, such as ROTEX and ETS-VII. In ten years, we expect to see more dexterous robots, such as the Space Dexterous Robotic Manipulator (SPDM), that can perform routine tasks such as changing out components under teleoperation. With intense effort, robots may be able to autonomously access and change-out obstructed components. Breakthroughs are needed to achieve advanced, autonomous troubleshooting and repair of arbitrary faults.

Surface and In-Space Human Assistance

Surface human EVA assistance robotic concepts are being explored by the EVA Robotic Assistant. In field tests with suited astronauts, it demonstrated the ability to follow

them while carrying tools, and to help them deploy a solar panel and cables. The space shuttle and space station remote manipulators have been used to move crew members from one location to another and to assist in moving assembly components. The teleoperated robots Robonaut and Ranger have demonstrated tasks such as handing over tools, holding objects for astronauts and shining lights on the ground. In ten years we expect autonomous robots to work in physical proximity to EVA crew members with very limited physical interaction. With intense effort, robots may be able to approach being limited teammates, with natural language and gesture interfaces and strong physical interaction. Arbitrary human level interaction requires breakthroughs.

Surface Mobility

Mobility is achieved through the interaction of many robotic capabilities to achieve safe and effective navigation in an environment. Complexity increases dramatically with the degree of autonomy employed. With limited autonomy: localizing in the environment, navigating while avoiding obstacles and collecting scientific information have been accomplished. Current flight demonstrated surface mobility is the 1997 Sojourner rover. Its capabilities are surpassed by the Mars Exploration Rovers destined for Mars in 2003, Figure 2.




Figure 2. Mars Exploration Rover

To achieve the longer durations and distances, greater science return, and reduced operations effort envisioned for future missions, enhanced robotic capabilities and increased robot autonomy are necessary. Significant capabilities include monitoring system state and health, acting in a resource-efficient manner, building maps, seeking targets of opportunity, and exploring to discover the unknown. Mechanical capabilities as well as energy and thermal issues are also relevant. These individual capabilities aggregate into the overall performance that

can be achieved in terms of duration, distance, speed, complexity, and reliability. Terrestrial robots Hyperion, Dante, and FIDO have demonstrated long-distance autonomous navigation, extreme terrain mobility, and relevant science operations, respectively. Simultaneous localization and mapping is largely solved in theory with remaining problems and methods for data association being advanced in coming years. Planning systems from terrain navigation to mission resource scheduling are today functional with sophistication and effectiveness that is expected to improve through the decade.

Surface Instrument Deployment

Sojourner was able to place forgoing instruments in contact with rocks from several meters away under supervised teleoperation over three or more command cycles. Terrestrial robots like K9 have demonstrated autonomous instrument placement. Rocky 7 and FIDO have demonstrated autonomous approaches to targets using visual means. In ten years, such systems will demonstrate sufficient robustness for deployment on missions. Intense effort is needed to deal with more complex situations, such as difficult terrain, occlusions and operations in highly confined areas.

Mission Planning and Sequence Generation

Current ground science planning tools allow planning with concurrency, flexible temporal conditions, and resource constraints with task-level, prioritized science input to generate verified sequence (example: Europa + MAPGEN). Scientists can work directly with the planning tools to generate a sequence of actions likely to be accepted by the flight engineering team. In ten years scientists may have full and direct control of the terrestrial rovers.

Onboard Science Planning and Perception

For terrestrial systems, the current state-of-art consists of onboard rover planners that maintain prioritized lists of science goals with multiple constraints between them, enabling fully autonomous operations for short durations (hours) in relatively simple outdoor environments (such as Antarctica). In ten years we expect steady improvements in robustness allowing fully autonomous operations for up to a day in environments similar to a able to seek patterns and anomalies and generate discovery plans to thereupon collect interesting scientific data at dramatically reduced operational effort. Performance at the level of a human scientist in the field is and will continue to be a major challenge. Without significant breakthroughs, the best systems will perform well within narrowly defined areas of expertise (as expert

systems do), but lack the general cognitive and perceptual abilities of a field scientist.

Additional Considerations

Robustness and interacting with robots at the *mission level* are two of several cross-cutting significant challenges that emerge in space robotics. Robustness is a challenge because robots must interact with complex environments, which may not be amenable to standard approaches to verification and validation. Furthermore, human level adaptability remains beyond the technological grasp of robotics. Robots that are autonomous and self-reliant-- able to address any fault through self-diagnosis and repair/recovery, and long-lived (years of operation) against the physical challenges of power, temperature, wear, and stability-- will remain a technological challenge. Interacting with robots at the mission level implies interpreting ambiguous instructions that the robot can only resolve through intimate knowledge of both the task and humans with which it has to interact.

We next turn to the case of human-robotic teams, and how they are best developed and utilized.

Human and Robotic Systems

As discussed above, we may expect that there will be some dramatic advances in the capabilities of planetary surface rovers in the next two decades. The question arises then as to whether such powerful robots will obviate the need for humans in exploring other planets -- the Moon and Mars in particular. This may well be the case for lunar exploration because the light distance of the Moon from Earth is only about a second and travel time to and from the Moon (e.g. for iterative sample return missions) is measured only in days. For Mars however, the round trip light distance ranges from about 10 minutes to about 40 minutes and one-way travel times approach a year so the exploration of Mars using robots controlled from Earth (including the return of samples to terrestrial laboratories) will be a slow business (as it has proven for the last 35 years).

The exploration of Mars using only humans is similarly problematic in that the martian surface is comparable in area to the terrestrial continents and, for a long time, human explorers will be confined in their in-person exploration to a region less than 1000 km across. So, the global exploration of this highly diverse planet is something that will surely need a combination of humans and robots. Even for their local in-person exploration astronauts will greatly benefit from the support of robots - since the number of astronauts will be limited and it will

be unsafe to conduct EVAs without at least a robotic "buddy".

We have begun to examine experimentally the question of how capable a robot controlled by a human at light distances of seconds (rather than 10s of minutes) can hope to be. We wish to know if "telepresence" is a potential reality or a chimera? In our pilot experiment we necessarily were required to *simulate* a mobile surface robot because, obviously, we expect major advances between now and ~2020 when the first human Mars exploration mission might plausibly be undertaken. As shown in Figure 3, the simulation consisted of a human-operated ATV equipped with several imaging instruments, ruggedized computers and a field communication system. We expect that by 2020 mobile rovers will be able to respond to high level commands in terms of navigation and data gathering -- in our simulation this was carried out by the ATV operator who received instructions via e-mail.



Figure 3. 2020-Class Rover Simulated With A Human-Operated ATV

The pilot field experiment was carried out in July 2002 at a Mars analog site, specifically the Houghton Crater on Devon Island in the Canadian Arctic, a location where logistical support including satellite communications have been developed over several years. A two person science team (both were geologists) operated the sim-rover from the NASA Ames' *FutureFlight Central*, shown in Figure 4, where panoramic images can be back-projected on large screens that occupy the 360° circumference of the facility. The situational awareness that this provides is remarkable and was a significant factor in motivating the experiment. The facility also had many individual monitors for the convenient display of higher resolution (than that of the panoramas) images of the far field, close-up images of the near field and macro images of individual samples.




Figure 4. Future Flight Central, Configured For Sim-Rover Control At A Distant Analog Site.

Several different sites outside the rim of Haughton Crater (not familiar to the science team) were selected for remote exploration with the goal of comparing the team's success in analyzing the site with evaluations that were carried out by geologists at Haughton Crater. Subsequent visits and traverses to these same locations were performed for comparison by spacesuited (Figure 5) and by unencumbered geologists.



Figure 5. Field Geologist Studying An Ejecta Block Near Haughton Crater While In The Hamilton-Sundstrand Prototype Suit.

Specific hypotheses to be evaluated were:

- In combination with high resolution orbital maps, digital imaging from a surface site can, under the right circumstances, enable an experienced, remote geologist to successfully determine the *ground truth* for the site in question
- A principal difference in terms of science productivity (i.e. time required to achieve an equal level of scientific insight) among the three cases (given unlimited communication bandwidth) is *latency* (which includes several factors and is irreducible beyond two-way light distance: zero for the in-person case, ~ 8 seconds for a

scientist/controller located at a Mars-Sun Libration Point, and tens of minutes for a scientist/controller on Earth).

(It should be noted that present robotic missions have a latency of about a day because the return of data depends on the infrequent availability of relay satellites. By 2020 it may be anticipated that direct, relatively high bandwidth communications from the Mars surface to Earth will be available and that by then the system latency could, in principle, be measured in tens of minutes.)

Communications from the Haughton base camp to NASA Ames were provided by satellite link and land line. The available bandwidth was expected to be 768 kbps. The number of individual nodes through which the data had to pass from the ATV to the FFC screens was more than twenty. Software "scripts" were written to control the acquisition of images by the camera systems and their implementation revealed that the system latency would be measured in minutes, thereby setting limits in what degree of telepresence the experiment could achieve.

Table 1. HORSE Camera Systems

Camera	Resolution	Image File Size	Script Time	Image Data
PanoCam	7363 X 704	2.3 to 4.3 MB	2:30	8 bits / RGB
Close-up Nikon	2560 X 1920	500 to 600 KB	0:05	8 bits / RGB
Kodak	2160 X 1440	1 MB	1:00	8 bits / RGB
High Resolution PanoCam	14726 X 1024	3.2 to 5.2 MB	4:58	8 bits / RGB / "High Res." Pano Cam setting
Backup Nikon	2560 X 1920	500 KB	0:05	8 bits / RGB

The table above demonstrates that the available digital imaging systems operate rather slowly, with typical

latencies of a few minutes, to which must be added the transmission time over a satellite link at ~ 700 kbps. Thus at this stage the light time delay was not a significant factor in the overall system latency.

The pilot experiment was a valuable learning experience and served to illuminate a number of issues that had not been obvious during the planning. First among these was the difficulty of navigating the ATV from one location where the science team was gathering data to their next desired study location. Part of the problem was the result of differences in coordinate systems on aerial images and on maps (a problem that was subsequently overcome) and part was the result of loss of situational awareness in moving from one location to another. We conclude that continuity of situational awareness is essential and we will investigate the use of frequent low resolution webcam imaging in future studies.

A number of factors appear to have led the science team to an overly rigid "flight plan" approach to their exploration, one which resembled a speeded up robotic mission approach rather than the slowed-down in-person approach that had been anticipated. The several factors in question were:

- Preoccupation with navigation issues
- Lack of continuous imaging
- Minimal preparation
- Overcast weather and limited resolution/contrast panoramas
- Inadequate time to learn and adjust to the experimental conditions
- A robotic mission mind-set

A comparison of the scientific insight of the remote team in comparison with that of the in-person exploration results remains to be carried out but it is already clear that this attempt to create a telepresence capability falls far short of what is needed and of what could be achieved with more effort.

All of these factors are ones that can be addressed before a follow-on experiment is undertaken. Plans are being made to carry out a more systematic experiment in conjunction with colleagues at the Johnson Space Center, this time in the desert South West where clear weather conditions can be reliably predicted and where the remote science team can travel to the sites and re-explore them to identify the various clues that we expect they will have missed during their remote effort. Contributing to the expected failure of a remote science team to fully characterize a field site will be the following:

- imaging coverage, resolution, stereography
- continuity of imaging

- system latency
- image display approaches
- instrumentation (besides imaging) on robot
- human-robot communication interface
- robotic autonomy
- end to end information system

The nature of the missed clues will allow us to judge what improvements can reasonably be made to the remote exploration technology and what is the likely limit in the effectiveness of telepresence operations.

In the next section, autonomous control concepts being developed at NASA are discussed. These concepts are being developed and demonstrated within the context of a broad range of mission scenarios, from single spacecraft to teams of spacecraft and rovers that cooperate in the exploration tasks.

Autonomous Control Concepts & Systems

Autonomous Control Concepts

The ability to autonomously monitor and control complex devices, such as a robotic explorer system, is critical to NASA's ability to accomplish many of its long-term space exploration goals. From the beginning of the space program in the late 1950's, control of spacecraft and systems have been managed by a large number of highly trained ground control personnel. This has its roots in the limited capability and massive size of computers of that early period. In addition, sensor data was telemetered to the ground, where a room full of experts would monitor each individual system's health and send commands to the spacecraft, either directly or via an astronaut. Over the past forty years there has been a radical shift in this paradigm, resulting largely from the advent of advanced computer technology. Automation has eased the burden of the ground controller and the astronaut, but often the tasks performed by the software are still quite rudimentary. This is a result of both computational resource limitations and the difficulty encountered when trying to develop, test, and validate software that provides the required functionality. As missions extend outward in the solar system, beyond the Earth-Moon system, the physical and fiscal realities of space exploration will require new control technologies.

Conceptually, the task of controlling a device such as a planetary rover is principally one of maintaining the system in a stable state while commanding transition of the device through a series of configurations designed to accomplish a sequence of goals in some optimal fashion. This task however, is often quite difficult due to the normal variations that occur within both the process and

the environment, limited observability into the current state of the device, and the potential of abrupt failures and degraded component performance. Traditionally, these

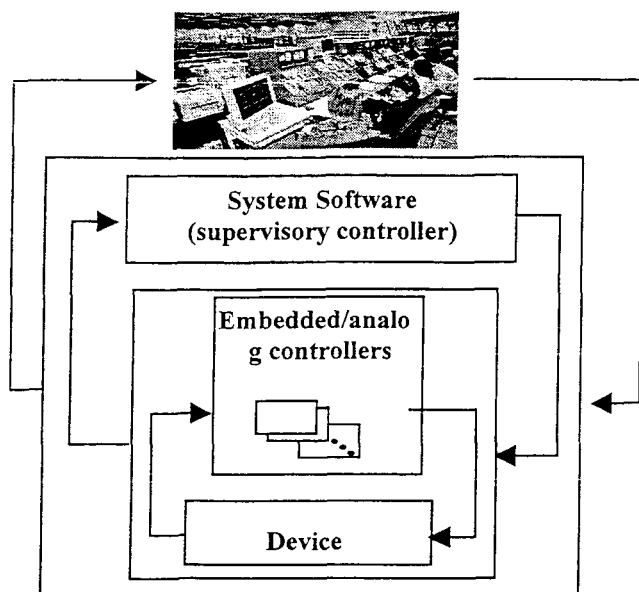


Figure 6. Tiered control architecture

problems have been solved through the use of a tiered architecture comprised of three levels, as shown in Figure 6 above:

1. **analog and embedded feedback controllers** to perform low-level regulatory functions,
2. **higher-level system software** to perform nominal command sequencing, and threshold monitoring to detect and respond to off-nominal conditions, and
3. **humans** to generate the command sequences, monitor the state of the device to detect off-nominal conditions, diagnose failures when they occur and select recovery actions in response to these failures.

While the capabilities provided by the system-level software have substantially increased over the past 40 years, the complexity of the missions undertaken by NASA have also increased. As a result, the requirements levied on the ground control team have increased, thus requiring larger ground support teams. This paradigm begins to break down for future missions, due both to greatly increasing time delays in communication as well as increasing costs of maintaining a larger ground support team. As a result, the role traditionally performed by the ground support team is being shifted to the system-level software, thereby greatly increasing the functionality required of this component.

Currently, the system-level software is developed by engineers who use commonsense understanding of the hardware and mission goals to produce computer code and control sequences that will allow a spacecraft, or other system, to achieve a particular goal while allowing for some (usually very small) amount of uncertainty in the environment. In developing this code, the engineer must reason through complex sub-system interactions to generate procedural code that can account for all the different combinations of failures and off-nominal conditions that might occur. As the functionality that is required of the system-level software increases, development, test, validation and maintenance of this software, using this traditional approach, becomes very difficult, if not impossible, as a result of the myriad of off-nominal conditions that the software is expected to handle. Furthermore, as the engineers gain a better understanding of how the device is behaving after deployment, it is often quite difficult to update the code to reflect the additional information that has been obtained.

Artificial Intelligence and Autonomous Control

As attempts are made to automate the processes that are traditionally performed by humans when monitoring and controlling these devices, it is clear that it will often be necessary to replicate the sophisticated inferencing capabilities exhibited by humans when performing this task. Over the last 4 decades, the field of artificial intelligence has been developing a variety of automated techniques that emulate a human's reasoning ability [3]. While the systems developed are far from performing at the visionary level exhibited by the HAL 9000 computer in *2001: A Space Odyssey*, such accomplishments as the victory of Deep Blue over Kasparov have demonstrated that sophisticated inferencing tasks can, indeed, be automated.

As an example of a first major step in providing an autonomous operating capability to an actual spacecraft, NASA developed and demonstrated, in flight, the Remote Agent (RA) autonomous control architecture. The next section will briefly describe that architecture and the subsequent flight experiment that led to a new space mission capability.

Remote Agent and the Deep Space 1 Mission

The Remote Agent architecture, developed collaboratively between NASA Ames Research Center (ARC) and the Jet Propulsion Lab (JPL), was part of the Deep Space One mission within the NASA New Millennium Program (NMP). It combined high-level planning and scheduling, robust multi-threaded execution, and model-based fault detection isolation and recovery,

into an integrated architecture that was able to robustly control a spacecraft over long periods of time [4-9]. The overall architecture of the Remote Agent, along with the additional elements incorporated for the Deep Space 1 mission, is diagrammed in Figure 7, and will be explained below.

Remote Agent itself was made up of three major components, each of which played a significant, integral role in controlling the spacecraft: Planner and Scheduler (PS)—produces flexible plans, specifying the basic activities that must take place in order to accomplish the mission goals. Smart Executive (EXEC)—carries out the planned activities. Mode Identification and Recovery (MIR)—monitors the health of the spacecraft and attempts to correct any problems that occur.

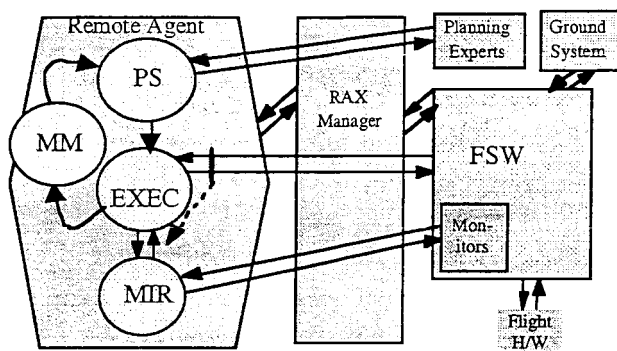


Figure 7. Remote Agent Architecture

These three parts worked together and communicated with each other to make sure the spacecraft accomplished the goals of the mission: EXEC requested a plan from PS. PS produced a plan for a given time period based on the general mission goals and the current state of the spacecraft. EXEC received the plan from PS, filled in the details of the plans, (eg., determined what spacecraft system actions must take place to complete the planned activities), and commanded the spacecraft systems to take the necessary actions. MIR constantly monitored the state of the spacecraft. It identified failures and suggested recovery actions. EXEC executed the recovery action or requested a new plan from PS that would take into account the failure. The parts of Remote Agent were constantly communicating (using inter-process communication) with each other and with the external components of the spacecraft. MIR received information regarding the state of different components from monitors located throughout the spacecraft. PS must receive information from planning experts in order to generate the plan. For example, the navigation system reports to PS regarding the spacecraft's current position, and the attitude-control system tells PS how long it will take to turn the spacecraft to a new position. Finally, EXEC

sends commands to other pieces of flight software that, in turn, control the spacecraft's systems or flight hardware.

As proof of the concept and capabilities, the Remote Agent software operated NASA's Deep Space 1 spacecraft and its futuristic ion engine during two experiments in May of 1999. For two days Remote Agent ran on the on-board computer of Deep Space 1, more than 60,000,000 miles (96,500,000 kilometers) from Earth. The tests were a successful step toward robotic explorers of the 21st century that will be less costly, more capable, and more independent from control from Earth.

Future exploration missions envision utilizing multiple rovers that are able to interact among themselves and thus greatly increase the science data return. The next section describes some of the techniques and technologies that will provide that capability.

Distributed Autonomous Systems

From the Terrestrial Planet Finder to robots helping each other scale cliffs on Mars, many future NASA mission concepts involve teams of tightly coordinated spacecraft/rovers in dynamic, partially understood environments. In order to maintain team coherence, each spacecraft must robustly respond to team coordination anomalies as well as local events. Currently manual techniques for implementing such teams are extremely difficult. These techniques involve either having one spacecraft tightly control the others or giving each spacecraft a separate activity sequence with explicit communication actions to coordinate with other spacecraft. While both approaches can handle two or three simple spacecraft, neither scales to larger populations or more complex spacecraft. New techniques are needed to facilitate managing populations of 3 or more complex spacecraft.

In general, autonomous spacecraft and rovers must balance long-term and short-term considerations. They must perform purposeful activities that ensure long-term science and engineering goals are achieved and ensure that they each maintain positive resource margins. This requires planning in advance to avoid a series of shortsighted decisions that can lead to failure. However, they must also respond in a timely fashion to a dynamic and partially understood environment. In terms of high-level, goal-oriented activity, the spacecraft must modify their collective plans in the event of fortuitous events such as detecting scientific opportunities like a sub-storm onset in Earth's magnetosphere or a Martian hydrothermal vent, and setbacks such as a spacecraft losing attitude control.

Whether they are orbiters, probes or rovers, coordinating multiple distributed agents introduces unique challenges. Issues arise concerning interfaces between agents, communication bandwidth, group command and control, and onboard capabilities. For example, consider a mission with a cluster of satellites simultaneously observing a point on a planet from different angles with different sensors. A certain level of communication capabilities will need to be assigned to each, possibly limiting the amount of information that can be shared between the satellites (and a ground station). The onboard capabilities also need consideration, including computing power and onboard data storage capacity. This will limit the level of autonomy each of the satellites can have. Finally, these issues also apply to multiple rover missions. A group of rovers might want to simultaneously measure vibrations caused by an explosion to determine the subsurface geology of an area on Mars. Developments are currently underway in team planning and execution to address these issues [10].

Team Plans

Current missions control multiple spacecraft by either giving one spacecraft a control sequence and having it treat the others as if they were virtually connected or giving each spacecraft its own command sequence with inserted communications actions for coordination between spacecraft. Neither approach scales as populations increase nor as members become more capable. While the first requires too much bandwidth, the second suffers from instabilities that lead to a variety of coordination failures.

These limitations can be overcome with hierarchical 'team plans' [11], which provide a rich representation for coordinated actions and communication while allowing the flexibility needed to adapt those plans in response to execution-time uncertainties. The representation for reactive team plans assumes a model of teamwork implemented on agents (like spacecraft) whose architectures explicitly support representing and reasoning about team goals, team plans, and team states. Team plans are hierarchical plans generalized to include team operators. Our approach builds on hierarchical reactive plans to provide a representation for explicitly defining constellation behavior. Giving each spacecraft the full constellation plan has several advantages over breaking up a plan into different role plans for each spacecraft. It enables a spacecraft's awareness of how its current activity relates to other spacecraft's activities. This information makes a constellation more robust by letting spacecraft monitor each other's progress. Also, this information facilitates autonomous recovery when unexpected events happen. When one spacecraft fails to

make progress, the related spacecraft can coordinate their responses.

Onboard Team Planning

While team plans provide a powerful way to specify team behaviors, teams operating in uncertain environments must be able to adapt those plans in response to unexpected events and opportunities. This capability is provided by onboard team planning, which extends single-spacecraft onboard planning to address issues of coordination and limited inter-agent communication bandwidth.

In order to facilitate this form of goal-based constellation commanding we are developing two complementary techniques. *Goal Distribution Planning* will let a designated lead spacecraft/rover plan with an abstract model of all followers in order to assign goals across the population. *Contract Networks* will let any spacecraft/rover serve as an auctioneer to distribute goals. These two approaches are actually two points in a spectrum of approaches where the leader gives its followers progressively more autonomy in deciding who satisfies which goals and how to satisfy a goal. While the first approach uses one spacecraft to collect all constellation information and generate a team plan/schedule, the second reduces communication overhead by spreading planning and scheduling activities across all spacecraft.

In addition to extending these approaches to work with continual planners, this research explores the spectrum by developing a hybrid system that uses both approaches.

Goal Distribution Planning (Master/Slave)

In the master/slave approach, Figure 8, one lead spacecraft embodies all four of the components and teleoperates the others. Our approach extends the MADS master/slave approach toward teamwork. In this approach a distinguished lead spacecraft will collect goals, generate a fleet plan, and broadcast it for execution. Just like a self-commanding spacecraft, the lead constellation spacecraft must respond in a timely fashion in a dynamic partially understood environment.

While this approach benefits from conceptual simplicity, it relies on an assumption that the leader's hardware proxies can continuously monitor the slaves' hardware, and this relies on high-bandwidth highly reliable communications. Reducing the requirements involves localizing reactive feedback loops by putting hardware proxies on all spacecraft, but this requires replicating the executive/diagnostician to appropriately manage the local

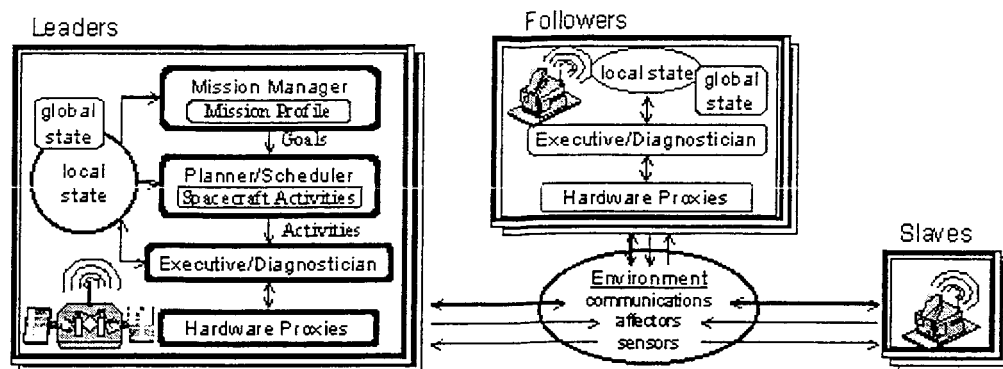


Figure 8. Software anatomy of leaders, followers, and slaves in an autonomous constellation.

hardware proxies. The final result is a constellation control architecture where the constellation's leaders, followers, and slaves depend on mission requirements.

Concluding Remarks

New capabilities for space exploration are becoming available that will significantly extend our reach and grasp. New robotic capabilities in key phases of of in-space assembly, planetary mobility and navigation, as well as scientific instrument placement will greatly increase our ability to explore and understand our exploration targets. With these new robotic capabilities will also come the ability to effectively team humans and robots to further extend our exploration goals. And, where the target of exploration exceeds the range of human presence, new, autonomous systems will provide the ability to explore and provide the knowledge and understanding that otherwise would be unavailable.

Acknowledgements

Portions of the work described in this paper were performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Particular thanks to Benjamin D. Smith, Manager of the Autonomy Technology Program at the Jet Propulsion Laboratory.

References

1. NEXT Robotic Survey 2002.
2. S. J. Hoffman and D. I. Kaplan, editors. *Human Exploration of Mars: The Reference Mission of the NASA Mars Exploration Study Team*. NASA Special Publication 6107. July 1997.
3. S. Russel and P. Norvig *Artificial Intelligence: A Modern Approach*, MIT Press 1995.
4. D. E. Bernard et Al. Design of the Remote Agent Experiment for Spacecraft Autonomy. *Proceedings of IEEE Aero-98*.
5. J. de Kleer and B. C. Williams, Diagnosing Multiple Faults, *Artificial Intelligence*, Vol 32, Number 1, 1987.
6. J. de Kleer and B. C. Williams, Diagnosis With Behavioral Modes, *Proceedings of IJCAI-89*, 1989.
7. J. de Kleer and B. C. Williams, *Artificial Intelligence*, Volume 51, Elsevier, 1991.
8. B. C. Williams and P. P. Nayak. Immobile Robots: AI in the New Millennium. In *AI Magazine*, Fall 1996.
9. N. Muscettola. HSTS: Integrating planning and scheduling. In Mark Fox and Monte Zweben, editors, *Intelligent Scheduling*. Morgan Kaufmann, 1994.